NASnews

September – October 1997 • Volume 2, Number 26

In This Issue...

DAO-NAS Partnership Steps Up With Two More SGI Systems

NASA's Data Assimilation Office (DAO), in collaboration with NAS Facility staff and Silicon Graphics Inc. (SGI), is migrating the Goddard Earth Observing System Data Assimilation System (GEOS-DAS) software from its current home at Goddard Space Flight Center to a permanent home at the NAS Facility. The DAO's function within NASA is to advance the state of the art in data assimilation--the process of turning observations into a model--and to provide assimilated datasets for the Mission to Planet Earth enterprise. To run the climate model and the assimilation algorithm, the office requires high-speed computing resources like those used at the NAS Facility.

To The Article...

NASnews

► Main Menu September – October 1997 • Volume 2, Number 26

DAO-NAS Partnership Steps Up With Two More SGI Systems

by Ayse Sercan

NASA's Data Assimilation Office (DAO), in collaboration with NAS Facility staff and Silicon Graphics Inc. (SGI), is migrating the Goddard Earth Observing System Data Assimilation System (GEOS-DAS) software from its current home at Goddard Space Flight Center to a permanent home at the NAS Facility. The DAO's function within NASA is to advance the state of the art in data assimilation -- the process of turning observations into a model -- and to provide assimilated datasets for the Mission to Planet Earth enterprise. To run the climate model and the assimilation algorithm, the office requires high-speed computing resources like those used at the NAS Facility.

With the arrival of two 64-processor Silicon Graphics Inc. Origin2000 systems in June, the collaborative efforts of NASA's Data Assimilation Office and the NAS Systems Division have stepped up considerably. The two organizations will maintain the systems in a partnership, with the NAS Facility providing the computing infrastructure and support, and the DAO providing cutting-edge hardware.

System Names Honor Pfaendtner

The software being migrated from Goddard will run on three Origin2000s -- the newest systems from SGI -- at the NAS Facility. In June, the DAO purchased and installed two new Origin2000s (jimpf0 and jimpf1), both 64-processor systems that will run the GEOS-DAS software. These systems -- collectively called jimpf -- are named after James W. Pfaendtner, chief of the Production Management Branch of the National Weather Service, who died unexpectedly in April 1995. Pfaendtner was a founding member of the DAO.

The first system (turing, also a 64-processor Origin2000), owned jointly by the DAO and the NAS Systems Division, is used for development and testing. It was installed at the NAS Facility in April (see NAS News, May-June '97).

Two-tiered Environment

In order to get near-production performance from these platforms, the project team has created a two-tiered environment. The plan is to use turing as a testbed, while jimpf will receive operating system and middleware updates from turing as they become stable. Mary Hultquist, Origin2000 project lead, explained that NAS will be doing development work on turing, and that "because the systems are identical in hardware and software, we can take any work we've done on turing and, once it's stable, move it onto DAO systems."

In addition to maintaining a separate development platform, the group is taking a number of precautions with the DAO systems. "We're going to keep a 'hot backup' of the last operating system version so that we can reverse any changes that cause problems in a short time. If a problem comes up that somehow wasn't noticed on turing, we can always go back to the stable environment quickly," said Jose Zero, parallel systems engineer, who is coordinating the software aspects of the migration. Also, jimpf0 and jimpf1 are independent systems. "They don't share resources internally, so if one goes down, you still have the second one," he said.

Need For Strong Change Control

One of the challenges of such an ambitious project is controlling all the changes made to the system. Hultquist explained that the group is keeping system backups of multiple steps in the process "so if anything goes wrong, we can fall back to the previous version without a problem. All changes to the systems are logged, and available via the Web."

Zero added: "This is a very large project with a lot of people working on it. Right now we've got a very strict configuration management system for the computing environment. All changes are logged and tested before they're released -- even to turing. The description of all the changes has to be coordinated by a single person, so there is one hundred percent coordination of all changes in the system."

The machines are being built, configured, and maintained by Archemedes de Guzman in the NAS parallel systems group, so there is close interaction in their development and maintenance. Zero is coordinating the DAS software changes, Hultquist coordinates system changes.

First Use in Production Environment

"We're the first group in the country to use these cutting-edge machines in a production-like environment," Zero noted. This environment will allow the DAO to better support NASA missions such as POLARIS (Photochemistry of Ozone Loss in the Arctic Region In Summer), SONEX (SASS Ozone and Nitrogen Oxide Experiment), and STRAT (Stratospheric Tracers of Atmospheric Transport). Currently, those missions are being supported by eagle, the ACSF CRAY C90, using the older GEOS-DAS software that requires fewer floating-point operations per second and less memory than the newer GEOS-DAS software.

Work on the project began in April with the acquisition of turing, and activity has stepped up since the arrival of Jimpf in June. The current focus of the GEOS-DAS project is the migration of its multitasking code and development of MPI (Message Passing Interface) code from Cray's Fortran 77 to a more generic version of Fortran 77, which will run on the Origin2000s. The goal is to complete the migration of the multitasking code by the end of June 1998, and to have the MPI version of the code ready to be used in December 1998.

The current work is the first in the collaboration between the NAS Systems Division and the DAO in the DAO's Earth-Observing System (EOS) project. According to the agreement between the two organizations, the DAO will provide funds for hardware and support, while NAS will provide a stable environment to run the GEOS-DAS software.

Ongoing Collaborations Benefit All

The DAO will be able to take advantage of the NAS computing expertise and infrastructure, as well as the disaster recovery capabilities of redundant systems. "Even though the machines are owned by the DAO, it's a partnership with NAS. If a NAS machine fails, the NAS [user] community can b riefly take advantage of DAO machines, and vice versa," Zero said. "The DAO is not just a customer, but a partner." He added that the DAO plans to expand this capability throughout the next year.

Both the DAO and the NAS Systems Division have collaborative relationships with SGI involving the Origin2000s. Each site has a memorandum of understanding in place with SGI for assistance with the development of key codes and for early software releases.

NASnews
▶ Main Menu
September – October 1997 • Volume 2, Number 26

Information Power Grid Will Give Scientists More Time for Research, Opportunities to Share Data

by Ayse Sercan

Managing access to computing resources is complex and consumes much of the time that could be spent doing research. In an ideal world, researchers would spend no time at all deciding which system to use. In addition, the way computing resources are managed today makes sharing of data difficult. Disparate computing resources keep disciplines stratified, so researchers often end up wasting time by unknowingly doing redundant work. And computing resources are often wasted because they are not allocated ideally -- a researcher who decides when and where to run a job is often not aware of the loads and priorities of all systems.

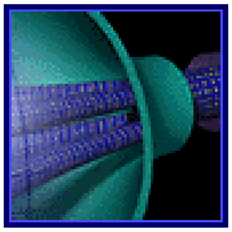

In order to address these issues, and taking a broader look at assisting in NASA's research and development missions, the NAS Systems Division is designing an "Information Power Grid" -- a twenty-year project that aims to seamlessly integrate computing systems, data storage, specialized networks, and sophisticated analysis software.

Like an electrical power grid, the Information Power Grid (IPG) will offer a steady, reliable source of computing power. Its scheduling software will dynamically and intelligently allocate resources among dispersed computing centers as they are needed.

Design Will Exploit Data Sharing

In the IPG model, each NASA center can purchase and maintain its own computing resources, investing in cutting-edge technology. At the same time, this design will help other centers share those resources. The IPG will automatically find the right system for the job, with allocations based on the workload content rather than site-specific capabilities. In addition, the IPG will be a data repository, allowing researchers to share data, re-manipulating it as needed. Ultimately, the IPG is meant to be not just for researchers using NASA systems, but a prototype for the implementation of a national computing matrix. The benefits of developing a system as complex as the IPG are not only the cost savings afforded by intelligently allocated resources. It will also provide assistance in collaboration, and allow a higher level of complexity in problems.

A New Approach to Research

The IPG will move the focus of researchers' interactions with the computer back to the research. Instead of asking, "What will the computers I have access to let me do?" they can simply ask, "What problem do I want to solve?" In addition, the IPG will allow interdisciplinary teams of researchers to work with the same data, possibly influencing each other's results.

Projects that are expected to benefit the most from the IPG are those in which interdisciplinary data is key -- such as air travel. Marisa Chancellor, former deputy division chief and one of the original IPG designers, predicted that the IPG could (for example) help cut aircraft accident rates and improve fuel efficiency by allowing the cross-referencing of earth-science weather databases with flight plans.

NASA will benefit directly from the IPG in a reduced cost of space launches, weather predictions for space launches, and analysis of space science data.

Behind The Scenes

The IPG will come with an intelligent interface, a knowledge-based environment that will help researchers prepare their problems for the computer, and will offer guidance on complex computing tasks.

Just behind the intelligent interface will be an application layer, where user applications and systems analysis tools run. In addition, there will be sophisticated data exploitation tools that allow researchers to analyze and manage their data. Hidden away from the users, the cluster/network operating system connects user applications to distributed resources. Process management and scheduling software will schedule and reserve processors, network services, and storage nodes as needed. The user will access these resources without needing to know exactly where they are or how they work.

The IPG will have a complex hardware back end, with computing nodes that include testbeds, computing prototypes, communications devices such as the routers and switches on NASA's Information Systems Network (NISN), and storage devices.

> 'It's one thing to say we're going to connect all these machines so they look as if they were one big computer. It's another thing to actually implement that.'

The IPG will be both a challenge and a boon to systems administrators. The challenge is in tying together distributed, heterogeneous systems, making them secure and reliable. But because the IPG is designed to be dynamically reconfigurable, computing resources can be added or deleted from the operational configuration as needed.

Challenge in 'Little Things'

Bill Nitzberg, NAS parallel systems group lead, has been working with Bruce Blaylock, chief engineer at NAS, to flesh out the technical aspects of the project and build near-term implementation plans. According to Nitzberg, "This project is really deceptive. The challenge is in the tiny little things, and the tiny little things, in this context, aren't so small any more. It's one thing to say 'we're going to connect all these machines so they look as if they were one big computer.' It's another thing to actually implement that."

Nonetheless, Nitzberg emphasized that even though standardization of hardware and software would make the project easier in the short term, "there are no plans to standardize hardware or homogenize NASA centers into one giant computing facility."

"That would be a bad move technologically," Nitzberg continued. "What we want to do with the Information Power Grid is take advantage of new, different hardware bought by other centers. It's harder to integrate those systems, but we need to stay on top of developments."

Big Picture is 'Solid,' Details in Progress

The Information Power Grid was conceived by a team comprised of NASA program managers who fund NAS Systems Division work, as well as NAS senior management.

More concrete details about requirements were added by a program development team, made up of NAS Facility customers in fields as disparate as integrated design systems, aviation operations systems, space science large-scale information management and simulation, space science operations, and autonomous systems.

While the big picture is fairly solid, the nuts and bolts of the system are still being discussed. A task force headed by Blaylock is currently working out short-term steps -- one-year and five-year plans.

Jim McCabe, a member of the NAS wide are networking team, explained that "there are three major problems to be worked out with the Power Grid: deciding what the user needs to see, tying all the resources together seamlessly, and managing all the information." The details of those problems are significant enough that the project is expected to take about twenty years to implement fully.

But while the committee is taking small steps towards their goal, "the little baby steps are going to get us there," McCabe explained, "It's one thing to be a visionary about the future, but without taking the small steps on the way there, you'll never get there."

A Tentative Timeline

According to McCabe, similar projects have gone wrong by making too many long-term -- and not enough short-term -- goals. McCabe stressed that the Information Power Grid is not just a pie-in-the-sky vision. "We're not just saying we're thinking about this, we're doing prototyping." Chancellor added that the IPG is more of a framework for the various projects the NAS Systems Division will be doing, a way of looking at the big picture rather than a radical departure from what is currently being done.

"There's the official Information Power Grid program, and then there's the work we're going to be doing anyway to make the NAS Facility more useful to our users. The power grid is a nice way to describe the overall vision, and a way to bring together other divisions, but even without the power grid as a framework, the projects will be done," she said.

The IPG will begin to be available to NAS Facility users around late 1998, when the first part of the tentative one-year plan goes into operation. Nitzberg said that the first appearance of the IPG will probably be some connectivity between Langley and Ames Research Centers, because the two sites already have a working relationship in the Metacenter project.

The IPG plan will undergo formal review before year end, at which time goals and milestones will be available. For more information, a "hotlist" with links to related topics and documents is located on the Web.

NASnews

▶ Main Menu September – October 1997 • Volume 2, Number 26

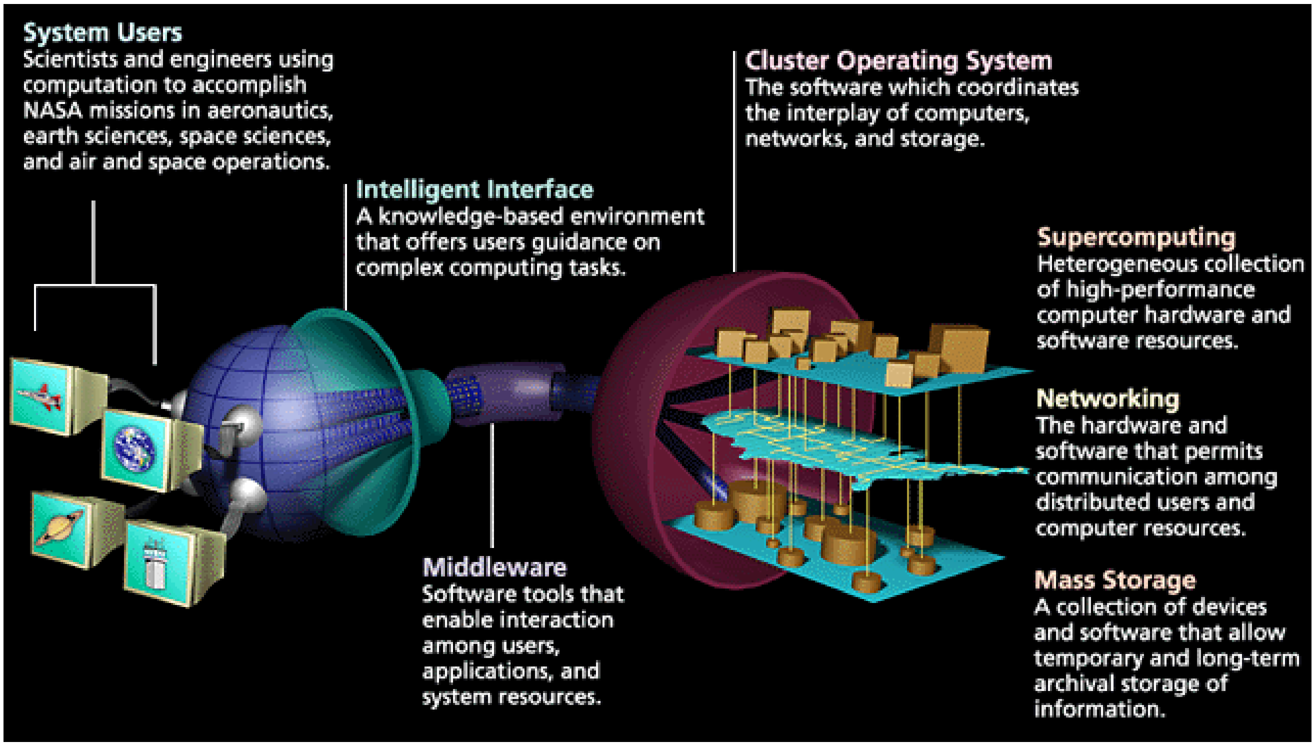

The aim of the Information Power Grid is to provide a reliable and seamless source of computing power. The challenges: providing an intelligent interface that is powerful but not time-consuming; coordinating an amalgamation of computing resources; and making the results available to researchers in a wide array of relevant disciplines. While implementation will span 20 years, results will be visible in the short term. The first goal: increased connectivity between Langley and Ames Research Centers, expanding on the working relationship of the Metacenter project.

Return to Main Article

NASnews
▶ Main Menu
September – October 1997 • Volume 2, Number 26

Software Corrects Distortion in Virtual Environment

by David Kenwright

A chance meeting between NASA Ames researchers Dan Delgado and Sam Uselton led to a surprisingly fast solution to a distortion problem that has plagued virtual reality systems for nearly a decade. Uselton, data analysis group lead in the NAS Systems Division, made a connection between earlier work on interacive particle tracing and Delgado's work in the Ames' Human Systems Technologies Branch.

One of the projects that is currently being conducted in that branch's Advanced Displays and Spatial Perception (ADSP) Laboratory involves correction of position and orientation measurements used in the creation of virtual environments. Researchers are testing a system for tracking head-mounted displays, which are used for several visualization and assembly tasks. One such task is the assembly of electrical wiring harnesses used in aircraft manufacturing (see graphic, below).

Accurate readings of position and orientation are needed to simulate precis ion assembly tasks in virtual environments. Electromagnetic sensors placed on the head-mounted display and the hand-held probe are used to obtain this information. Here, an engineer assembles a prototype electrical wiring harness f or an aircraft. The wiring "blueprint" appears on a the head-mounted display. Graphic courtesy of Boeing Airplane Corp.

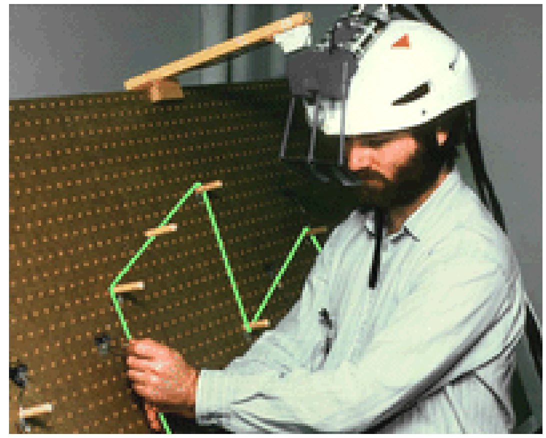

The ADSP lab uses commercial electromagnetic position measurement systems to find the orientation and location of one to four sensors within a region that is about three meters from a transmitter. The transmitter emits a magnetic field that the sensors receive and analyze. Generally, one sensor is attached to an operator's head-mounted display and another to a hand-held probe. This combination of hardware allows engineers to design and assemble components in a virtual environment without the expense of building and storing a physical prototype.

Magnetic Field Warps Virtual Environment

Electromagnetic trackers are designed to operate in a metal-free environment. In the real world, most buildings -- including the ADSP Lab -- have steel components in the floor, walls, and furniture. This disturbs the magnetic field produced by the transmitter and causes distortion in the virtual environment. The distortion gets worse as the operator moves further away from the transmitter, and can be as high as 20 percent (or about one foot) at the perimeter. From the operator's point of view, this distortion makes the world look warped.

"It's like living in distorted space," explained Stephen Ellis, group leader for the ADSP Lab, "and this causes problems such as motion sickness." A more serious problem, Ellis pointed out, is that operators have difficulty moving their hands in a straight line -- a critical task in constructing wiring harnesses.

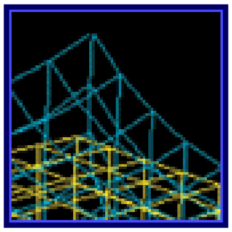

To identify the distortion, Ellis's group (including B. D. Adelstein, Stefan Baumeler, Hans Jense, and Mark Young) measured positions and orientations from the electromagnetic sensor at regular intervals throughout the room. The image at left was constructed from their data. Without any distortion, the image would be a regular cubic or Cartesian mesh.

The problem then became how to convert the distorted readings of the tracker to the correct Cartesian coordinates, and do this in real time (60-120 Hertz) on an inexpensive workstation.

Connection Leads to Collaboration

When Uselton heard about this problem from Delgado, he immediately realized it was similar to the point location problems that are routinely solved with CFD visualization software. Point location is the task of finding out exactly where you are in a curvilinear mesh. As part of the calculations, the earlier NAS-developed software performs a transformation into an orthogonal coordinate system -- and that component was precisely what the ADSP team needed.

Dramatic Results in Days

Within a week, Delgado and colleague Rick Jacoby modified and incorporated the software originally developed for particle tracing into the ADSP lab's system to correct the distortion -- with dramatic results. The virtual environment no longer looked warped and the errors in tracker position were reduced to millimeters.

In July, the ADSP team completed controlled experiments with novice operators in both the distorted and undistorted virtual environments. The role of the head tracker in these experiments was to test human effectiveness in performing simple tests like following the path of a moving object in the virtual environment (see Quicktime video, below). Analysis of the data is currently underway, but early results indicate that user training is more effective in the undistorted environment.

Example of the 3D tracking tasks used to evaluate the performance impact of removal of tracker distortions.

Latency = 0 Frames, Quicktime video, 160x120, 1.4 MB

Latency = 6 Frames, Quicktime video, 160x120, 1.2 MB

According to Ellis, the distortion correction software will be useful to anyone using an electromagnetic position tracking system. Commercial organizations outside Ames have already expressed interest in the project.

For more information on NAS's point location software, contact David Kenwright.

David Kenwright works in the NAS data analysis group, and developed the point location software that corrected the distortion in the Ames Advanced Displays and Spacial Perception lab's visual environment.

NASnews
Main Menu
September – October 1997 • Volume 2, Number 26

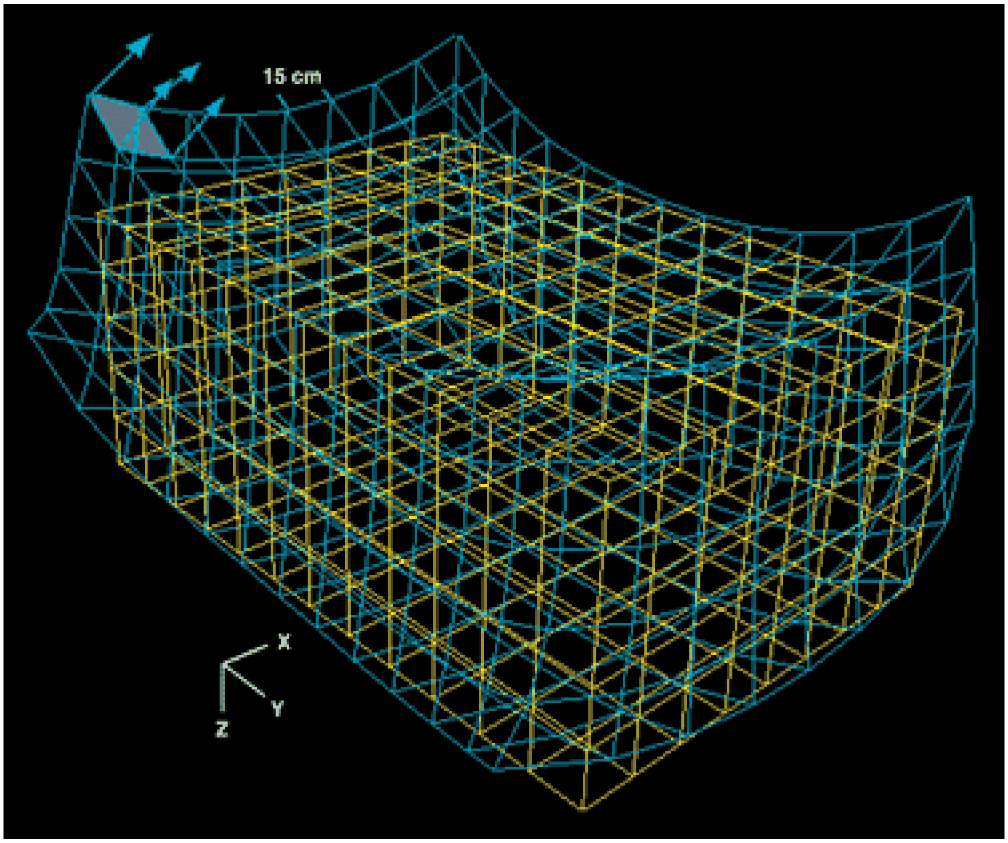

Position and orientation measurements were taken at regular intervals throughout the NASA Ames Advanced Displays and Spatial Perception (ADSP) Laboratory using a commercial electromagnetic tracker. Metal structures in the lab distorted the measurements, producing a "warped world." Using software developed by David Kenwright at the NAS Facility, this distortion can now be corrected in real time. This image shows the degree of before (blue) and after (yellow) position correction. Graphic courtesy of ADSP Lab.

Return to Main Article

NASnews
▶ Main Menu September – October 1997 • Volume 2, Number 26

Users Win by Utilizing $BIGDIR on the Crays

by Daniel DePauk

Despite numerous pleas from the NAS scientific consultants, many users are not taking advantage of the latest high-speed processing techniques. Because the Cray supercomputers at the NAS Facility continue to be in high demand, it is important that all jobs on these systems run as efficiently as possible. This article is the first in a series on how to improve job performance and turnaround time, reduce allocation charges, and increase overal system performance. Some of these methods are easy to incorporate‹others are not. Users are stronly encouraged to consider all that are feasible when creating new jobs.

This article presents an easy method for obtaining potentially impressive performance gains: using the $BIGDIR directory and reserving disk space with the Session Reservable File System (SRFS.)

The NAS Facility's high speed processor (HSP) staff randomly checks the HSP systems for overall performance problems. By observing the recent workload on the Cray systems (vonneumann and eagle, two CRAY C90s; and the four newton systems, CRAY J90s) the group has found that users continue to use their home file systems as current working directories, instead of reserving disk space and using $BIGDIR.

The following table shows the low percentage of user jobs that reserved

$BIGDIR space with SRFS during a five-day period in late July.

System	SRFS Jobs
vonneumann	32%
eagle	29%
newton	39%

A $BIGDIR directory is created for every log-in session (including batch jobs). Since space on the /big filesystem can be used without a reservation, an examination of total I/O on /big is appropriate. The amount of space being used and the I/O activity on the /big file system on vonneumman and eagle indicate that the percentage of jobs using $BIGDIR without a reservation is higher than those reserving disk space with SRFS.

On the newton systems, the majority of jobs using $BIGDIR also make a reservation with SRFS, resulting in fewer $BIGDIR jobs. (See examples of setting up scripts to reserve disk space with SRFS and using $BIGDIR. NAS News, May-June '97.)

Differences Between Home and $BIGDIR

The filesystem containing the $BIGDIR directories, called /big, performs the best on all of the HSP systems because it is directly attached to disks with high data transfer rates. Of these systems, vonneumann, eagle, and newton1 have home directories on filesystems that are directly attached to disks and network devices with *lower* data transfer rates. These home directories are better suited for file storage and interactive use, rather than for batch jobs. In addition, newton2, 3, and 4 access the home directories over a HIPPI network, which provides an even slower data transfer rate.

The home directories also have Cray's Data Migration Facility (DMF) implemented for the superhome file systems. DMF can cause spikes in I/O requests during migration activities, resulting in unpredictable performance deviations in programs using home directories.

The $BIGDIR directory is available to jobs without reserving space with SRFS. However, space on /big is only guaranteed with a SRFS reservation. Jobs can exceed the amount of space reserved with SRFS on $BIGDIR if non-reserved space is available on /big.

To reserve space on the filesystem for $BIGDIR, the PBS (Portable Batch System) parameter string at the beginning of the job is:

```
#PBS -1 srfs_big=<space_needed>
```

The correct amount of space needed for a job can be obtained in the Resource Utilization Summary provided at the end of every PBS job on the HSP systems.

$BIGDIR Benefit on newtons

Using $BIGDIR reduces the probability that a job will fail due to the loss of newton1. Jobs performing I/O to $BIGDIR are unaffected by the unavailability of the home directories. The reason: when newton1 becomes unavailable due to a system or network failure, the home directories become unavailable on the other newton systems. So, jobs running on newton2, 3 or 4 that perform I/O between the home directories will fail when newton1 is unavailable.

Improvements Demonstrated

To demonstrate the performance differences between writing to the home filesystem and /big, a test program was produced that writes 512-word buffers of real numbers, creating a 10-MW file. Random reads and writes were then performed on the file. These I/O operations transferred a total of 809 MW of data. The test program ran twice on vonneumann, eagle, and each of the newton mainframes on a normally loaded system. The first run used the home filesystems, the second used $BIGDIR.

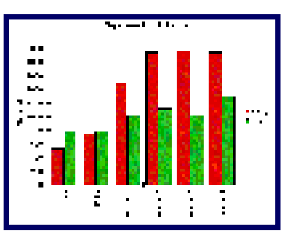

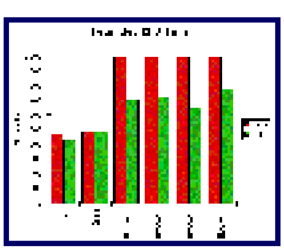

Using the Cray job accounting information command "ja" to obtain performance numbers, two graphs were produced. These graphs (left) directly address two concerns that many users have about their jobs: elapsed (wall-clock) time for a job and total job CPU time charged.

Because the CRAY C90 architecture of vonneumann and eagle can handle I/O more effectively, the differences in job CPU time are not as great between the home filesystem and $BIGDIR. Note that in both graphs, the difference is greater on the CRAY J90 architecture of newton1 than on either CRAY C90 system. There is an additional jump in both elapsed time and total job CPU time with the use of Network File System accessing users' homes filesystems on newton2, 3, and 4.

The percentage of saving depends on which system is being used, as well as the load on the system when the program is run. When the test programs ran, eagle's workload was just enough to prevent idle time, and few I/O-intensive jobs were running. The newton4 system does not have the I/O bandwidth to disk devices, causing the savings there to be less. The overall results do show that it is better to run using $BIGDIR than the home directories. The wall-clock time was reduced as much as 88 percent, with a 34 percent reduction in CPU time.

The Bottom Line

The benefits of investing a little time to modify your scripts to use $BIGDIR and SRFS are improved job turnaround, reduced CPU time, and increased job reliability. In addition, as more of you begin to use $BIGDIR, overall system performance will improve.

For more information or assistance with improving job performance, contact NAS User Services at (415) 604-4444 or (1-800) 331-8737, or send email to **nashelp@nas.nasa.gov**.

Daniel DePauk is a systems analyst in the high speed processor group. He has 18 years of experience in large-scale scientific computing, 11 of them at Ames Research Center. His primary focus is on system performance and tuning for all Cray systems at the NAS Facility.

NASnews
▶ Main Menu
September – October 1997 • Volume 2, Number 26

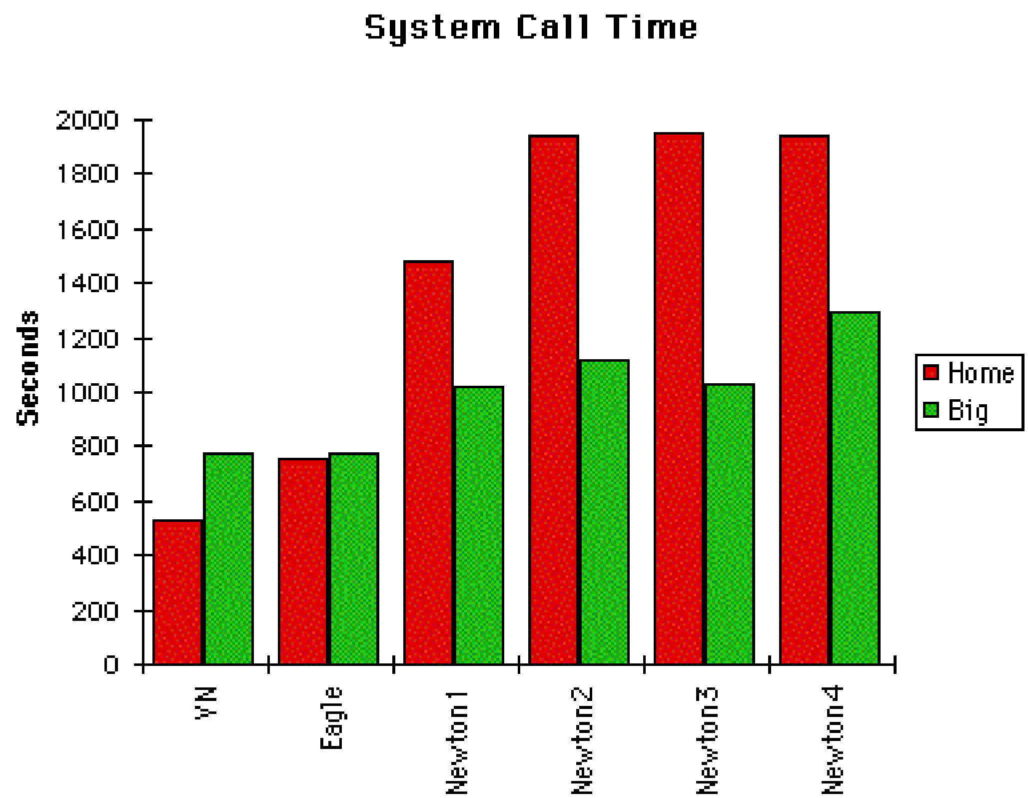

Elapsed (wall-clock) time for a job

Return to Main Article

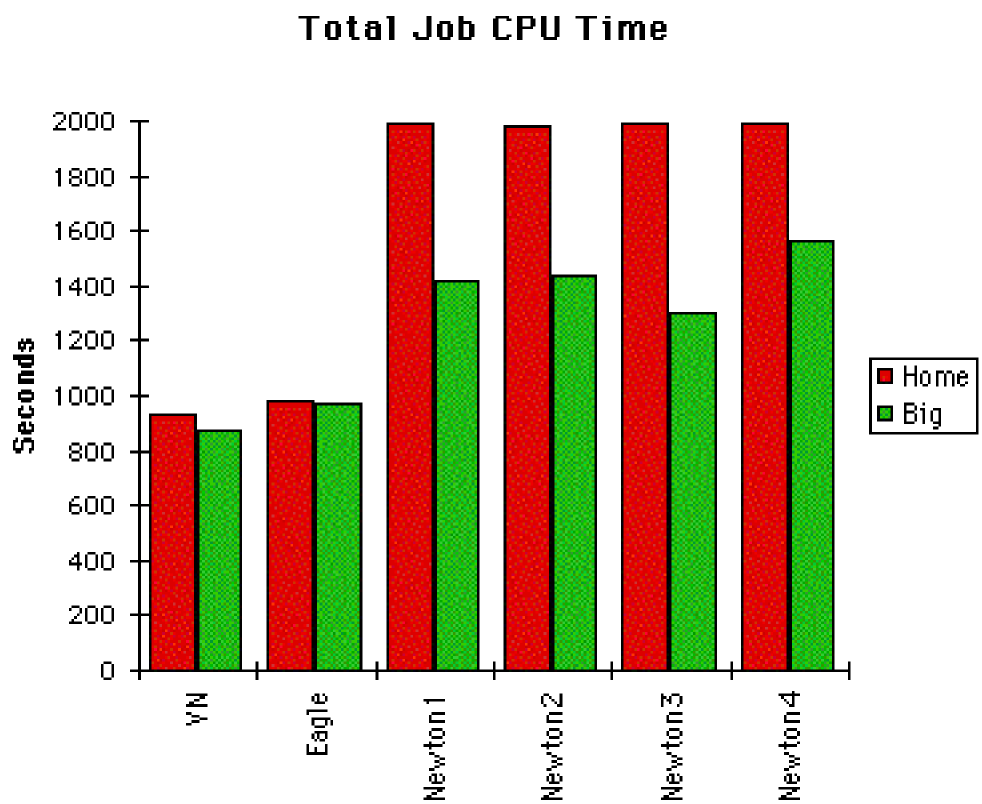

Total job CPU time charged

[Return to Main Article]

NASnews

▶ Main Menu September – October 1997 • Volume 2, Number 26

Parallel Graphics Library Lets Researchers Interact With Simulations

by Thomas W. Crockett

The traditional method for visualizing the output from supercomputer simulations has been to save snapshots of the data, then transfer them to a high-end graphics workstation for postprocessing and rendering. This approach has several drawbacks, including the need to move potentially massive datasets across the network, and difficulty in obtaining visual feedback while the application is in progress.

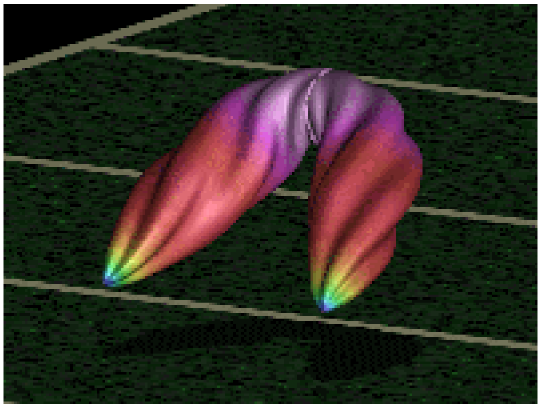

One frame from an animation sequence of a hairpin vortex rendered using the Parallel Graphics Library (PGL). Red and magenta depict regions of large angular velocity on the surface of the vortex tube, while flow in blue and green areas has lower angular velocity. The grooved indentations follow the path of streamlines on the tube's surface, providing a sense of the vortical motion. Graphic courtesy of Thomas W. Crockett.

click to see 400x300 gif of image, 66K

mpeg animation, 320x240, 1.55MB

An alternative, as described in the NAS Technical Summary " Parallel Graphics and Visualization Techniques for Distributed Memory Architectures,"; is to exploit the available processing power to perform the graphics and visualization operations in place, as the data is being created. With this approach, a parallel software renderer is coupled directly to the application code, replacing the hardware rendering engine found in graphics workstations. The output then becomes a stream of images, which can be compressed "on the fly" and transmitted to the user's workstation for real-time viewing. Because the supercomputer is performing most of the graphics work, inexpensive desktop workstations have sufficient power to decompress and display the images.

Images are often much more compact than the simulation data from which they are derived. For example, an uncompressed, full-color, full-screen image requires at most 3 to 4 megabytes (MB) of storage. The use of compression, color reduction, and lower-resolution images can often reduce the image data to a size that can be transmitted satisfactorily over local area networks at rates of several frames per second. For slower long-haul links, more aggressive lossy compression techniques can be employed to maintain adequate frame rates, with higher quality images being stored on the supercomputer's disk for later viewing.

Interactivity can be provided by means of a graphical user interface that runs on the user's workstation and sends rendering and visualization commands back to the parallel application. This allows the user to monitor the progress of a simulation, modify viewing and visualization parameters to explore the dataset, and potentially adjust critical parameters or terminate the job if it is not producing the desired results.

In order for this approach to be successful, efficient parallel algorithms are needed to perform the graphics and visualization operations. This has been an area of ongoing research for several years within NASA--sponsored in large part by the High Performance Computing and Communications' Computational Aerosciences Program and in the Department of Energy's national labs and the international academic community.

One result of this effort is a parallel 3D graphics library developed by researchers at the Institute for Computer Applications in Science (ICASE). The library, called PGL (Parallel Graphics Library), is available on the World Wide Web. In addition, the initial version of PGL is available on the NASA Metacenter IBM SP2s. The code is currently supported for SP2s and Paragons, but is designed to be portable to other message-passing parallel systems.

PGL is also available in sequential versions, and is currently being ported for use on networks of workstations. The workstation port of PGL will probably be tested on the davinci cluster at the NAS Facility.

The parallel rendering paradigm does have some limitations. For runtime use, display rates and interactivity are dependent on the quality of the network connection between the workstation and the parallel system. Congested long-haul routes or low-bandwidth connections may provide unsatisfactory performance. Of course, these delays also inhibit the downloading of large datasets for use with traditional postprocessing visualization techniques. It is expected that improvements in networking technology will eventually provide most research and development institutions with sufficient long-haul bandwidth to support image-intensive applications such as parallel rendering.

Beyond bandwidth requirements, parallel systems that support run-time visualization must be operated in a manner that facilitates interactive use, and should have adequate network connectivity to allow several users to generate image streams simultaneously without undue interference.

The NAS Technical Summaries, give overviews of this and other research projects conducted on NAS Facility computers. More information on the ICASE visualization and graphics research is also available.

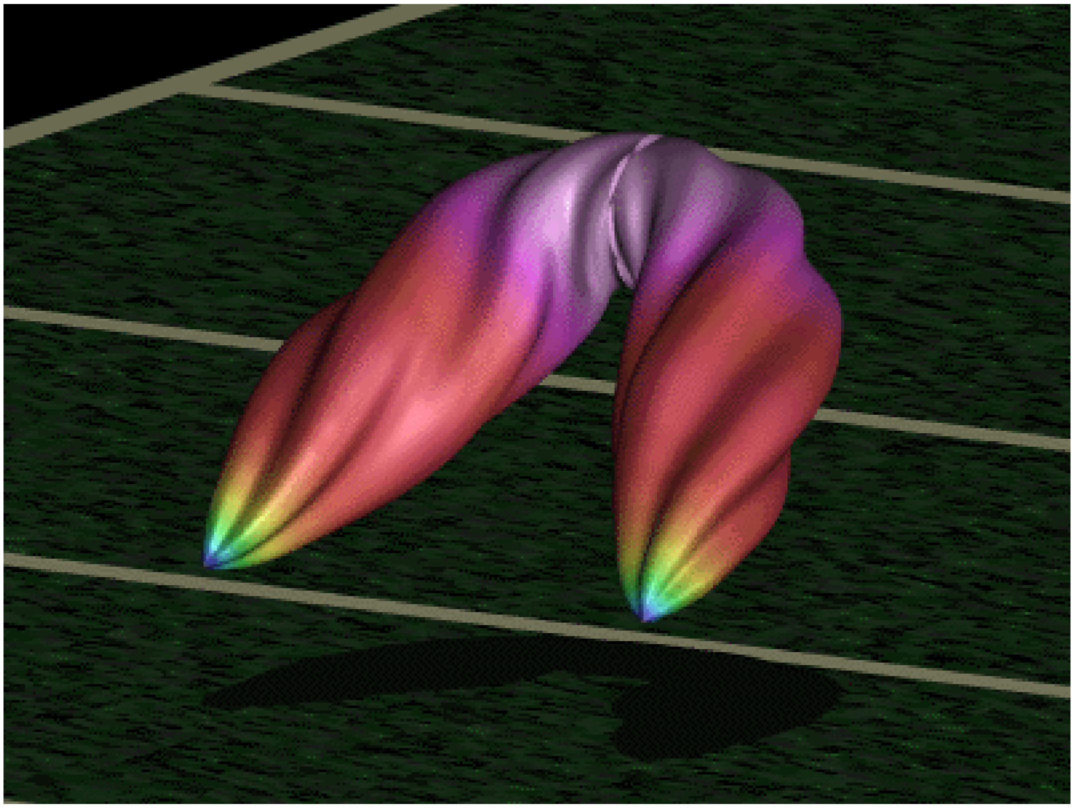

NASnews
▶ Main Menu
September – October 1997 • Volume 2, Number 26

NAS Researchers Present Work on Large Datasets, Unsteady Flows at Visualization Conference

by Ayse Sercan and Peter Adams

Four researchers from the NAS Systems Division will present their work at this year's IEEE Visualization '97 conference in Phoenix, October 19-24. The eighth annual conference will feature seminars and presentations on topics covering advances in visualization algorithms and techniques, emerging technologies, and developments in a diverse range of visualization applications, from archaeology to aerospace.

Paging Speeds Visualizations

Michael Cox and David Ellsworth will present their paper, "Application-Controlled Demand Paging for Out-of-Core Visualization," which describes their improvements to the Unsteady Flow Analysis Toolkit (UFAT), designed to speed visualizations of CFD datasets on workstations. Cox and Ellsworth have implemented changes in UFAT that allow it to divide its input datasets before loading them. These changes allow small workstations with limited memory to perform visualization tasks that were previously unmanageable.

"The way that data are usually visualized is that the entire dataset is copied into memory, and then it is visualized. But if you're talking about a 100-gigabyte dataset -- or even a 15-gigabyte dataset -- that's not possible. So what we do is take advantage of the fact that when you're visualizing your data, you don't need every single piece of data to get a visualization, by paging in only the data that are necessary," explained Cox.

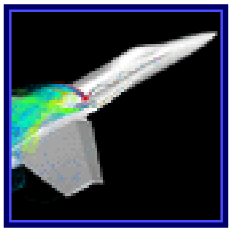

The original version of UFAT loads two timesteps of the dataset in memory during program execution. Cox and Ellsworth's experiments showed that their modification of UFAT was as good as or faster than the original when memory was ideal -- and outperformed the original by a wide margin in every situation where memory was lacking -- indicating that paging makes visualization of unsteady flows practical on rather small systems.

In addition, their tests showed that the most efficient method for storing multidimensional data is in "cubed" form, as opposed to the "flat" form with reference locality in only one dimension.

The pair plan to test additional techniques to reduce the memory required for very large datasets, and to integrate their work in the visualization group's Field Encapsulation Library (FEL).

Visualizing Vector Data in Unsteady Flow

Han-Wei Shen and David L. Kao will present their paper, "UFLIC: A Line Integral Convolution Algorithm for Visualizing Unsteady Flows," which outlines UFLIC, their new algorithm for visualizing vector data in unsteady flow fields.

The time-dependent algorithm designed by Shen and Kao uses Line Integral Convolution (LIC) as its underlying form. In order to model unsteady flows, it uses a successive feed-forward convolution method that maintains the coherence between animation frames and a time-accurate value depositing scheme that models local changes in the flow. The feed-forward method does not create each successive frame from the original, but creates it with the previous frame's output to ensure that consecutive frames are highly coherent. In addition, by using a high-pass filter, image clarity and texture contrast are maintained. Time-accurate value depositing ensures that as the algorithm repeatedly computes the evolution of the pathlines, the pixels involved advect only in a forward pathline direction -- reflecting physical reality, because particles in flows cannot advect backward in time.

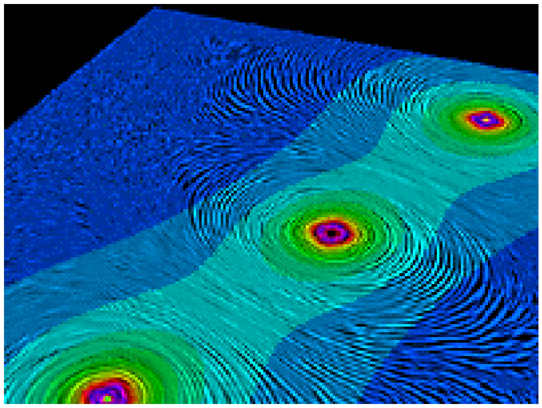

Snapshot in an animation sequence of a two-dimensional flow simulation with dynamic vortices, created with UFLIC, the Unsteady Flow Line Integral Convolution software. The two outer vortices orbit clockwise around the central vortex, which spirals counterclockwise. The flow pattern is color-mapped by velocity magnitude, which is very high near the vortex center (magenta), and slower at the edges (darker blue). The UFLIC algorithm can accurately capture the dynamic features of this unsteady flow simulation. The animation will be shown at the eighth annual Visualization '97 conference, to be held in Phoenix, October 19-24, as part of a presentation by developers Han-Wei Shen and David L. Kao, both in the NAS data analysis group. Graphic by Han-Wei Shen.

Shen and Kao found UFLIC to be very effective in capturing dynamic features in unsteady flow fields; however, some areas of the visualization need further work. For example, blurring and inconsistent line widths may appear in areas where the flow is changing rapidly in velocity or direction. Also, because the pattern is displayed through texture mapping, aliasing effects are visible at different resolutions.

Future plans include work on speeding up and parallelizing the algorithm, and dealing with issues of memory management.

More information about NAS research in scientific visualization is available on the NAS data analysis group web page.

NASnews
Main Menu
September – October 1997 • Volume 2, Number 26

In this issue...

NAS Researchers Present Work on Large Datasets, Unsteady Flows at Visualization Conference

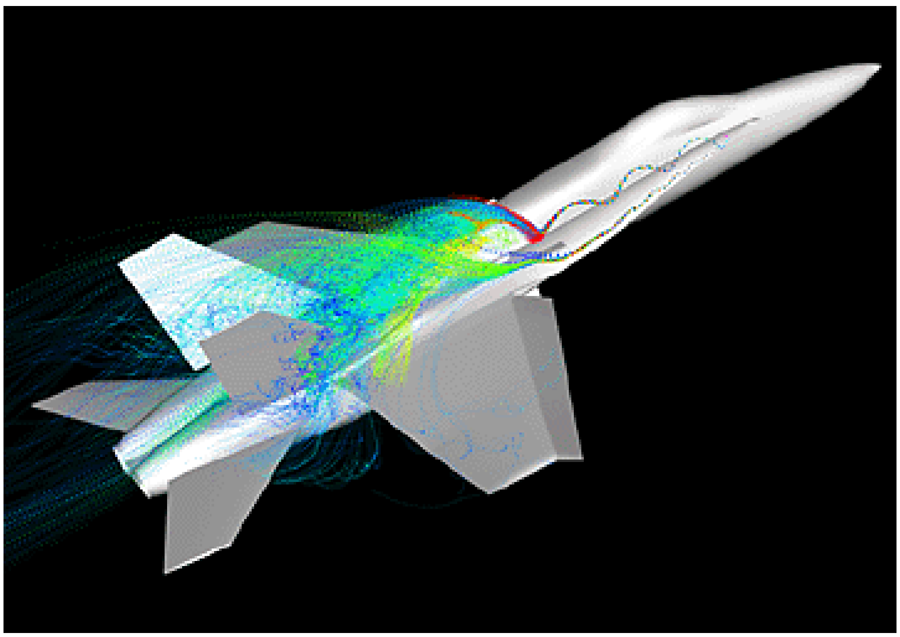

Michael Cox and David Ellsworth used this simulation of the vortex behavior of an F-18 in flight to test their adaptations to the NAS-developed Unsteady Flow Analysis Toolkit (UFAT), a set of tools for analyzing large, unsteady flow datasets.

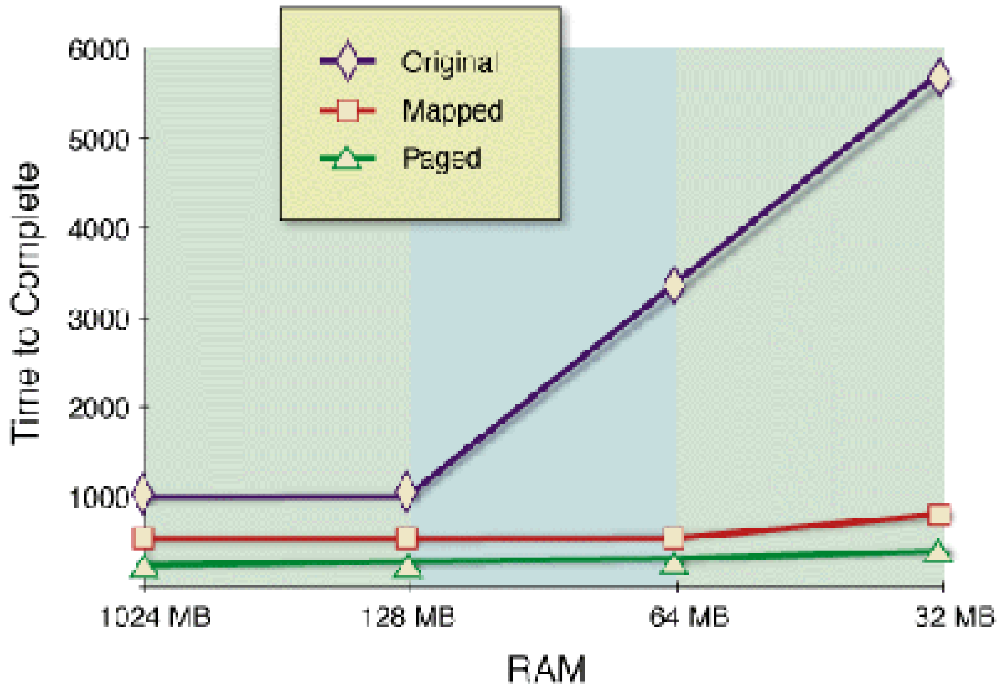

The chart above shows the clear improvement in performance the library offers. The original version of UFAT lags in low-memory situations. Cox and Ellsworth's paging system is clearly faster, with little change in performance as memory availability decreases, while a mapped version of the same toolkit is just slightly slower than paging.

Return to Main Article

NASnews

Main Menu September – October 1997 • Volume 2, Number 26

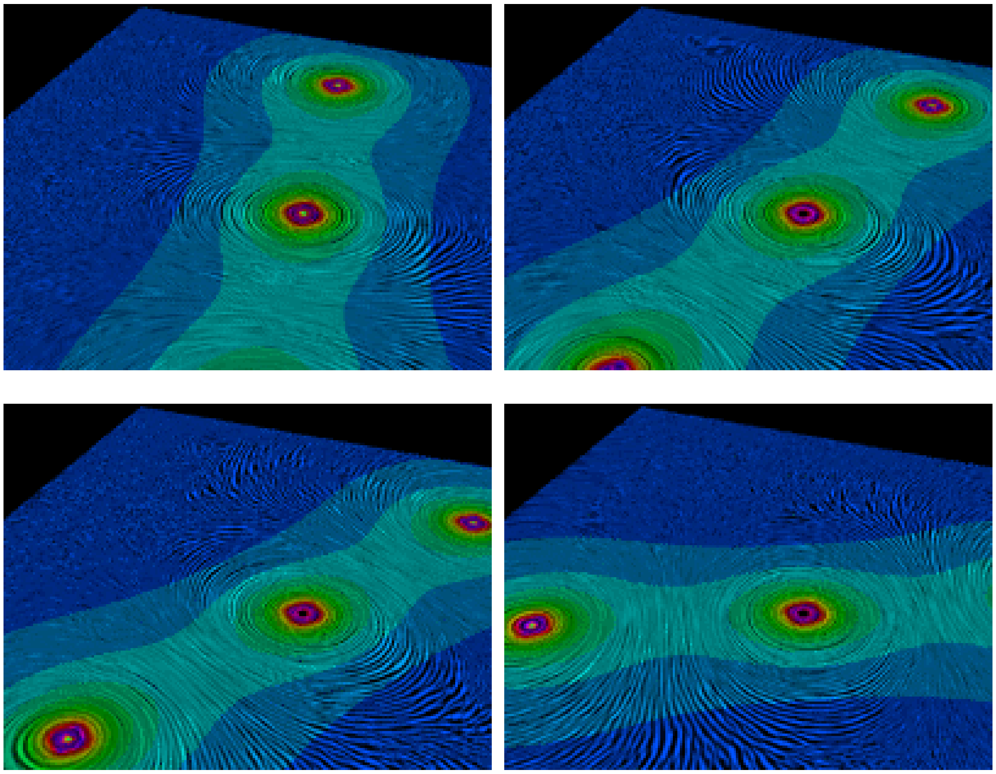

A time-sequence of a two-dimensional flow simulation with dynamic vortices. The simulation has two vortices orbiting clockwise around the central vortex, which spirals counter-clockwise, The flow pattern is color-mapped by the velocity magnitude. The UFLIC algorithm accurately captures the dynamic features in this unsteady flow simulation. Flow simulation dataset courtesy of Ravi Samtaney; image created by David Kao and Han-Wei Shen.

Return to Main Article

NASnews
Main Menu
September – October 1997 • Volume 2, Number 26

In This Issue...

Latest NAS Technical Reports Available Online

by Ayse Sercan

Here are summaries of NAS Technical Reports published in spring 1997 and available on the World Wide Web. To access specific reports, find the links corresponding to the technical report numbers indicated below.

"Communication Studies of DMP and SMP Machines," by Andrew Sohn and Rupak Biswas. An investigation of the interplay between programming paradigms and communication capabilities of parallel machines; in particular, the IBM SP2 distributed-memory multiprocessor and the SGI POWERCHALLENGE array symmetric multiprocessor. Experimental results indicate that the communication performance of the multiprocessors is consistent with the size of messages. The SP2 is sensitive to message size but yields much higher communication overlapping because of the communication coprocessor. The SGI array is not highly sensitive to message size and yields low communication overlapping. (NAS-97-005)

"Formation of Carbon Nanotube Based Gears: Quantum Chemistry and Molecular Dynamics Simulations of the Electrophilic Addition of o-Benzene to Fullerenes, Graphene, and Nanotubes," by Richard Jaffe, Jie Han, and Al Globus. The pathways for o-benzene addition to naphthalene, C_{60} and a model of curved polycyclic aromatic hydrocarbons (chosen to represent the carbon nanotube used in gear simulations) are considered. Density Functional Theory, a variant of *ab initio* quantum chemistry, molecular mechanics, and dynamics calculations are used for this study. The molecular mechanics calculations are carried out using the potential energy force field developed by Brenner. In addition, the Brenner potential energy formulation is assessed for possible use in the companion molecular dynamics simulations. (NAS-97-006)

"Tetrahedral and Hexahedral Mesh Adaptation for CFD Problems," by Rupak Biswas and Robert C. Strawn. Two unstructured mesh adaptation schemes for problems in CFD are presented. The procedures allow localized grid refinement and coarsening to efficiently capture aerodynamic flow features. The first procedure is for purely tetrahedral grids; unfortunately, repeated anisotropic adaptation may significantly deteriorate the quality of the mesh. Hexahedral elements, on the other hand, can be subdivided anisotropically without mesh quality problems. Computational results indicate that the hexahedral adaptation procedure is a viable alternative to adaptive tetrahedral schemes.
(NAS-97-007)

"Efficient Load Balancing and Data Remapping for Adaptive Grid Calculations," by Leonid Oliker and Rupak Biswas. A novel method to dynamically balance processor workloads, with a global view, showing the implementation and integration of all major components in the dynamic load balancing strategy for adaptive grid calculations. Mesh adaption, repartitioning, processor assignment, and remapping are critical components of the framework that must be accomplished rapidly and efficiently so as not to cause a significant overhead to the numerical simulation. The potential bottlenecks of mesh repartitioning and data remapping are resolved, and the framework remains viable on a large number of processors. (NAS-97-008)

"HARP: A Fast Spectral Partitioner," by Horst D. Simon, Andrew Sohn, and Rupak Biswas. The major factor that hindered the widespread use of partitioning unstructured graphs was long execution times. A new inertial spectral partitioner, called HARP, quickly partitions the meshes at run-time in a manner that works efficiently for real applications in the context of distributed-memory machines. A parallel MPI version of HARP has also been implemented on the IBM SP2 and the CRAY T3E. Parallel HARP, running on 64 processors, can partition a mesh containing more than 100,000 vertices into 64 subgrids in about half a second. These results indicate that graph partitioning can now be truly embedded in dynamically changing real-world applications. (NAS-97-009)

"Improved Boundary Conditions for Cell-centered Difference Schemes," by Rob F. Van der Wijngaart and Goetz H. Klopfer. Cell-centered finite-volume (CCFV) schemes have certain properties that make them attractive for the solution of compressible fluid flow equations. Unfortunately, they lead to slow convergence for numerical programs that utilize them. This paper presents one method for investigating and improving convergence of CCFV schemes, focusing on the effect of the numerical boundary conditions. The key to this method is the computation of the spectral radius of the iteration matrix of the entire discretized system of equations, not just of the interior point scheme or the boundary conditions. (NAS-97-011)

"Adaptive Load-Balancing Algorithms Using Symmetric Broadcast Networks," by Sajal K. Das, Daniel J. Harvey, and Rupak Biswas. Three novel, symmetric broadcast network (SBN)-based load-balancing algorithms are proposed, analyzed, and implemented on an IBM SP2. A thorough experimental study with Poisson-distributed synthetic loads demonstrates that these algorithms are very effective in balancing system load while minimizing processor idle time. They also compare favorably with other existing load-balancing techniques. Additional experiments performed with real data demonstrate that the SBN approach is effective in adaptive computational science and engineering applications where dynamic load balancing is extremely crucial. (NAS-97-014).

"Toroidal Single Wall Carbon Nanotubes in Fullerene Crop Circles," by Jie Han, investigates the energetics and structure of circular and polygonal single-wall carbon nanotubes (SWNTs) using large-scale molecular simulations on the IBM SP2 at the NAS Facility, motivated by their unusual electronic and magnetic properties. (NAS-97-015)

In This Issue...

Order NAS Technical Seminar Videos on the Web

Many research seminars presented at the NAS Facility are available through a videotape loan program. Information on past training events are available, as are procedures for obtaining training materials.

Here are summaries of the most widely attended technical seminars at the NAS Facility for the last several months.

"The Gradient Path Toward Design Tools for Nanotechnology Based on Electron Density," presented by Preston J. MacDougall, Middle Tennessee State University. Electron density is an observable property of matter, which can be calculated very accurately with state-of-the-art quantum chemistry software (see graphic, right). In position space, topological analysis has revealed the physical manifestations of concepts such as atoms, bonds, and electron pairs. This analysis is being used for molecular design purposes, such as pharmaceutical research. Density in momentum space can also be calculated with topological analysis. This provides model-independent information, which correlates well with known electron transport properties. The new methods and their evolution are discussed. (6/9/97)

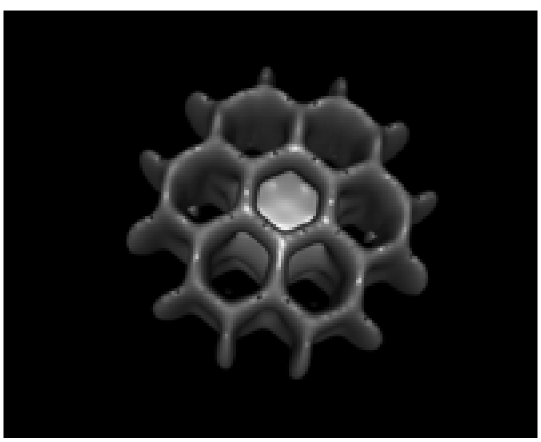

Preston MacDougall and Creon Levit developed a nano-design tool by interfacing the electronic structure program GAUSSIAN-94, with the Flow Analysis Software Toolkit (FAST). The surface shown is a momentum diagram for the electrons in a triangular array of six lithium atoms. The solid areas in this three-layered honeycomb indicate preferred vectors in the array. At the center of the surface are the slowest valence electrons, which are highly laminar in all directions. More distant surfaces indicate electrons at higher velocity. The "holes" in the honeycomb indicate directions in which electrons do not smoothly travel. Diagrams such as this will allow scientists a more intuitive understanding of the properties of the materials they are investigating.

"MP/Express--Distributed Visualization Software," presented by Anastasios Kotsikonas and Susan Margaret Ramsden. A technical overview of MP/Express. Following a discussion of project goals and how parallel modules are written, data distribution is discussed. This includes: descriptions of the LayoutSpec object components, which determine particular distribution patterns; the mapping table, which is fundamental to distribution functions; and the distribute and collect methods. Implementation of the "exchange neighbors" function is introduced, with brief descriptions of the inquiry functions. (6/26/97)

"Tools and Applications for High Performance Computing: R&D at the Vienna Center for Parallel Computing (VCPC)," presented by Barbara Chapman, Vienna Center Director. The VCPC was established in 1995 as a center for high-performance computing (HPC). The center collaborates with a number of organizations throughout Europe, focusing on the creation of mature applications and programming tools for parallel systems. Ongoing projects include the porting of two industrial codes under HPF. The discussion includes HPC and networking in a European framework, as well as the center's efforts in application parallelization and tool development. (7/8/97)

"Mars Direct," presented by Robert Zubrin. The United States has chosen a far-term view of manned missions to Mars, when in fact, all the necessary technology is already available. With an aggressive program based on small spacecraft launched directly to Mars by boosters, the US can reach Mars within a decade. The keys to success are traveling light and living off the land. (7/10/97)

"Performance Evaluation for Parallel Program Optimization," presented by Xian-He Sun, Computer Science Department, Louisiana State University. Parallel programming has been hampered by a lack of predictive performance evaluation (that is, a fast implementation may slow when the system or problem scales up). To promote scalability, four models of parallel speedup are given. Also introduced are the ideas of metric scalability and range comparison. The first measures the performance gain after a computing power increase, the latter compares the performance of programs over a range of ensemble and problem sizes, pinpointing the intersection between scale and performance. The development of a system, SCALA, is also discussed. Experimental results show that SCALA's range comparisons are feasible and effective for parallel computing. (7/15/97)

"Material Simulations at Multiple Length and Time Scales: Semiconductor Surfaces, Interfaces and Nanotubes," with Jerry Bernholc, Department of Physics, North Carolina State University. New theoretical methods and parallel computing power allows the study of several hundred independent atoms completely from first principles. While these simulations are truly predictive and informative, important limitations still remain. After briefly reviewing several recent advances, emphasis turns to the properties of nanotubes, where complex physical simulations necessitate the use of other techniques. (7/16/97)

"Multi-Disciplinary Computing Environment (MDICE) (Phase II SBIR)," presented by Gary Hufford, CFD Research Corporation. The MDICE system provides an environment for manipulating body multi-disciplinary applications requiring grid remeshing. Third party application programmers can operate in the environment without source code access. Grid remeshing is provided in the environment, supporting geometrical deformations with grid updating for a variety of schemes. Discussed topics included: a high level overview of MDICE; MDICE's Generalized Object Library; third party use of MDICE; the control module; and the group's future plans. (7/17/97)

"Conflict-Free Access to Data Structures and Distributed Computing," presented by Sajal K. Das, director of the Center for Research in Parallel and Distributed Computing, University of North Texas. In mapping data structures it is important that different subsets of nodes (templates) be accessible in parallel without memory conflict. To solve this problem, different forms that allow interconnection without interference must be analyzed. These forms include k-ary trees, binomial trees, and hypercubes. As applications, different template types play crucial roles in optimal implementation of different areas. The intricacy of this mapping led to the concept of oriented templates. The seminar ends with the presentation of open problems. (7/21/97)

NASnews
Main Menu
September – October 1997 • Volume 2, Number 26

Credits

Executive Editor: Thomas Lasinksi

Editor: Jill Dunbar

Staff Writer: Ayse Sercan

Contributing Writers: Peter Adams, Thomas W. Crockett, Daniel DePauk, David Kenwright

Image Coordinator: Fay Pattee

Online Page Layout and Graphics: Joel Antipuesto, Fay Pattee, Chris Gong, Rosemary Wadden

Other Contributors: Richard Anderson, David C. Banks, William J. Bene, Michael Cox, Dan Delgado, James Donald, Stephen Ellis, David Ellsworth, Clayton Guest, Mary Hultquist, James Jones, David L. Kao, Jim McCabe, George Myers, Bill Nitzberg, Marcia Redmond, Ricky Rood, Han-Wei Shen, Bart A. Singer, Jose Zero